arXiv:1408.5032v1 [stat.ML] 21 Aug 2014

On the Sample Complexity of Subspace Learning*

Alessandro Rudi, Guille D. Canas, Lorenzo Rosasco

August 22, 2014

Abstract

A large number of algorithms in machine learning, from principal component analysis (PCA), and its non-linear (kernel) extensions, to more recent spectral embedding and support estimation methods, rely on estimating a linear subspace from samples. In this paper we introduce a general formulation of this problem and derive novel learning error estimates. Our results rely on natural assumptions on the spectral properties of the covariance operator associated to the data distribution, and hold for a wide class of metrics between subspaces. As special cases, we discuss sharp error estimates for the reconstruction properties of PCA and spectral support estimation. Key to our analysis is an operator theoretic approach that has broad applicability to spectral learning methods.

1 Introduction

The subspace learning problem is that of finding the smallest linear space supporting data drawn from an unknown distribution. It is a classical problem in machine learning and statistics, with several established algorithms addressing it, most notably PCA and kernel PCA Jolliffe (2005); Schölkopf et al. (1997). It is also at the core of a number of spectral methods for data analysis, including spectral embedding methods, from classical multidimensional scaling (MDS) Borg and Groenen (2005); Williams (2002), to more recent manifold embedding methods Tenenbaum et al. (2000); Roweis and Saul (2000); Belkin and Niyogi (2003), and spectral methods for support estimation De Vito et al. (2010). Therefore knowledge of the speed of convergence of the subspace learning problem, with respect to the sample size, and the algorithms' parameters, is of considerable practical importance.

Given a measure ρ from which independent samples are drawn, we aim to estimate the smallest subspace S_ρ that contains the support of ρ. In some cases, the support may lie on, or close to, a subspace of lower dimension than the embedding space, and it may be of interest to learn such a subspace S_ρ in order to replace the original samples by their local encoding with respect to S_ρ.

While traditional methods, such as PCA and MDS, perform such subspace estimation in the data's original space, other, more recent manifold learning methods, such as isomap Tenenbaum et al. (2000), Hessian eigenmaps Donoho and Grimes (2003), maximum-variance unfolding Weinberger and Saul (2004, 2006); Sun et al. (2006), locally-linear embedding Roweis and Saul (2000); Saul and Roweis (2003), and Laplacian eigenmaps Belkin and Niyogi (2003) (but also kernel PCA Schölkopf et al. (1997)), begin by embedding the data in a *feature space*, in which subspace estimation is carried out. Indeed, as pointed out in Ham et al. (2004); Bengio et al. (2004b,a), the algorithms in this family have a common structure. They embed the data in a suitable Hilbert space $\mathcal{H}$, and compute a linear subspace that best approximates the embedded data. The local coordinates in this subspace then become the new representation space. Similar spectral techniques may also be used to estimate the support of the data itself, as discussed in De Vito et al. (2010).

*This paper is the extended version of (Rudi et al., 2013)

While the subspace estimates are derived from the available samples only, or their embedding, the learning problem is concerned with the quality of the computed subspace as an estimate of S_ρ (the true span of the support of ρ). In particular, it may be of interest to understand the quality of these estimates, as a function of the algorithm's parameters (typically the dimensionality of the estimated subspace).

We begin by defining the subspace learning problem (Sec. 2), in a sufficiently general way to encompass a number of well-known problems as special cases (Sec. 4). Our main technical contribution is a general learning rate for the subspace learning problem, which is then particularized to common instances of this problem (Sec. 3). Our proofs use novel tools from linear operator theory to obtain learning rates for the subspace learning problem which are significantly sharper than existing ones, under typical assumptions, but also cover a wider range of performance metrics. While the more technical parts of the proofs are postponed to the Appendices, a full sketch of the main proofs is given in Section 3.1, including a brief description of some of the novel tools developed. We conclude with experimental evidence, and discussion (Sec. 5 and 6).

2 Problem definition and notation

Given a measure ρ with support M in the unit ball of a separable Hilbert space $\mathcal{H}$, we consider in this work the problem of estimating, from n i.i.d. samples $X_n = \{x_i\}_{1 \le i \le n}$, the smallest linear subspace $S_\rho := \overline{\operatorname{span}(M)}$ that contains M.

The quality of an estimate $\hat{S}$ of S_ρ, for a given metric (or error criterion) d, is characterized in terms of probabilistic bounds of the form

$$\mathbb{P}\left[d(S_\rho, \hat{S}) \le \varepsilon(\delta, n, \rho)\right] \ge 1 - \delta, \quad 0 < \delta \le 1. \tag{1}$$

for some function ε of the problem's parameters. We derive in the sequel such high probability bounds.

In the remainder the metric projection operator onto a subspace S is denoted by P_S, where $P_S^2 = P_S^* = P_S$ (every P is idempotent and self-adjoint). We denote by $\|\cdot\|_{\mathcal{H}}$ the norm induced by the dot product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ in $\mathcal{H}$, and by $\|A\|_p := \sqrt[p]{\operatorname{Tr}(|A|^p)}$ the p-Schatten, or p-class norm of a linear bounded operator A Retherford (1993, p. 84).

2.1 Subspace estimates

Letting $C := \mathbb{E}_{x \sim \rho} x \otimes x$ be the (uncentered) covariance operator associated to ρ, it is $S_\rho = \overline{\operatorname{Ran} C}$ (see Proposition C.3). Similarly, given the empirical covariance $C_n := \frac{1}{n} \sum_{i=1}^n x \otimes x$, we define the *empirical subspace estimate*

$$\hat{S}_n := \operatorname{span}(X_n) = \operatorname{Ran} C_n$$

(note that the closure is not needed in this case because $\hat{S}_n$ is finite-dimensional), and the *k-truncated (kernel) PCA subspace estimate* $\hat{S}_n^k := \operatorname{Ran} C_n^k$, where C_n^k is obtained from C_n by keeping only its k top eigenvalues. Note that, since the PCA estimate $\hat{S}_n^k$ is spanned by the top k eigenvectors of C_n, then clearly $\hat{S}_n^k \subseteq \hat{S}_n^{k'}$ for $k < k'$, and therefore $\{\hat{S}_n^k\}_{k=1}^n$ forms a nested family of subspaces (all of which are contained in S_ρ).

As discussed in Section 4.1, since kernel-PCA reduces to regular PCA in a feature space Schölkopf et al. (1997) (and can be computed with knowledge of the kernel alone), the following discussion applies equally to kernel-PCA estimates, with the understanding that, in that case, S_ρ is the span of the support of ρ *in the feature space*.

2.2 Performance criteria

In order for a bound of the form of Equation (1) to be meaningful, a choice of performance criteria d must be made. We define the distance

$$d_{\alpha,p}(U,V) := \|(P_U - P_V)C^\alpha\|_p \tag{2}$$

between subspaces U, V, which, as shown in Proposition C.1 in the Appendix, is a metric over the space of subspaces contained in S_ρ, for $0 \leq \alpha \leq \frac{1}{2}$ and $1 \leq p \leq \infty$. Note that $d_{\alpha,p}$ depends on ρ through C but, in the interest of clarity, this dependence is omitted in the notation.

While of interest in its own right, it is also possible to express important performance criteria as particular cases of $d_{\alpha,p}$. In particular, the so-called **reconstruction error** Maurer and Pontil (2010):

$$d_R(S_\rho, \hat{S}) := \mathbb{E}_{x\sim\rho}\|P_{S_\rho}(x) - P_{\hat{S}}(x)\|^2_{\mathcal{H}}$$

is $d_R(S_\rho, \cdot) = d_{1/2,2}(S_\rho, \cdot)^2$ (see Proposition C.4).

Note that d_R is a natural criterion because a k-truncated PCA estimate minimizes a suitable error d_R over all subspaces of dimension k. Clearly, $d_R(S_\rho, \hat{S})$ vanishes whenever $\hat{S}$ contains S_ρ and, because the family $\{\hat{S}^k_n\}_{k=1}^n$ of PCA estimates is nested, then $d_R(S_\rho, \hat{S}^k_n)$ is non-increasing with k.

As shown in Maurer and Pontil (2010), a number of unsupervised learning algorithms, including (kernel) PCA, k-means, k-flats, sparse coding, and non-negative matrix factorization, can be written as a minimization of d_R over an algorithm-specific class of sets (e.g. over the set of linear subspaces of a fixed dimension in the case of PCA).

3 Main results

Our main technical contribution is a bound of the form of Eq. (1), for the k-truncated PCA estimate $\hat{S}^k_n$ (with the empirical estimate $\hat{S}_n := \hat{S}^n_n$ being a particular case), whose proof is postponed to 3.1.

We begin by bounding the distance $d_{\alpha,p}$ between S_ρ and the k-truncated PCA estimate $\hat{S}^k_n$, given a known covariance C.

Theorem 3.1. *Let $\{x_i\}_{1\leq i\leq n}$ be drawn i.i.d. according to a probability measure ρ supported on the unit ball of a separable Hilbert space $\mathcal{H}$, with covariance C. Assuming $n > 3$, $0 < \delta < 1$, $0 \leq \alpha \leq \frac{1}{2}$, $1 \leq p \leq \infty$, then the following holds uniformly for $k \in \{1, \ldots, n\}$:*

$$\mathbb{P}\left(\forall\, 1 \leq k \leq n \;:\; d_{\alpha,p}(S_\rho, \hat{S}^k_n) \leq 3t_k^\alpha \left\|C^\alpha(C + t_k I)^{-\alpha}\right\|_p\right) \geq 1 - \delta \tag{3}$$

where $t_k = \max\{\sigma_k, \frac{9}{n}\log\frac{n}{\delta}\}$, and σ_k is the k-th top eigenvalue of C.

We say that C has *eigenvalue decay rate of order r* if there are constants $q, Q > 0$ such that $qj^{-r} \leq \sigma_j \leq Qj^{-r}$, where σ_j are the (decreasingly ordered) eigenvalues of C, and $r > 1$. From Equation (2) it is clear that, in order for the subspace learning problem to be well-defined, it must be $\|C^\alpha\|_p < \infty$, or alternatively: $\alpha p > 1/r$. Note that this condition is always met for $p = \infty$, and also holds in the reconstruction error case ($\alpha = 1/2, p = 2$), for any decay rate $r > 1$.

Knowledge of an eigenvalue decay rate can be incorporated into Theorem 3.1 to obtain explicit learning rates, as follows.

Theorem 3.2 (Polynomial eigenvalue decay)**.** *Let C have eigenvalue decay rate of order r. Under the assumptions of Theorem 3.1, it is, uniformly in $k \in \{1, \ldots, n\}$, with probability $1 - \delta$*

$$d_{\alpha,p}(S_\rho, \hat{S}^k_n) \leq \begin{cases} Q' k^{-r\alpha+\frac{1}{p}} & \text{if } k < k^*_n \qquad \textit{(polynomial decay)} \\ Q' {k^*_n}^{-r\alpha+\frac{1}{p}} & \text{if } k \geq k^*_n \qquad \textit{(plateau)} \end{cases} \tag{4}$$

where it is $k^*_n = \left(\frac{qn}{9\log(n/\delta)}\right)^{1/r}$, and $Q' = 3\left(Q^{1/r}\Gamma(\alpha p - 1/r)\Gamma(1 + 1/r)/\Gamma(1/r)\right)^{1/p}$.

The above theorem guarantees a drop in $d_{\alpha,p}$ with increasing k, at a rate of $k^{-r\alpha+1/p}$, up to $k = k_n^*$, after which the bound remains constant. The estimated plateau threshold k^* is thus the value of truncation past which the upper bound does not improve. Note that, as described in Section 5, this performance drop and plateau behavior is observed in practice.

The proofs of Theorems 3.1 and 3.2 rely on recent non-commutative Bernstein-type inequalities on operators Bernstein (1946); Tropp (2012), and a novel analytical decomposition. Note that classical Bernstein inequalities in Hilbert spaces (e.g. Pinelis (1994)) could also be used instead of Tropp (2012). However, while this approach would simplify the analysis, it produces looser bounds, as described in Section 3.1.

If we consider an algorithm that produces, for each set of n samples, an estimate $\hat{S}_n^k$ with $k \geq k_n^*$ then, by plugging the definition of k_n^* into Eq. 4, we obtain an upper bound on $d_{\alpha,p}$ as a function of n.

Corollary 3.3. *Let C have eigenvalue decay rate of order r, and Q', k_n^* be as in Theorem 3.2. Let $\hat{S}_n^*$ be a truncated subspace estimate $\hat{S}_n^k$ with $k \geq k_n^*$. It is, with probability $1-\delta$,*

$$d_{\alpha,p}(S_\rho, \hat{S}_n^*) \leq Q' \left(\frac{9\left(\log n - \log \delta\right)}{qn} \right)^{\alpha - \frac{1}{rp}}$$

Remark 3.4. *Note that, by setting $k = n$, the above corollary also provides guarantees on the rate of convergence of the empirical estimate $S_n = span(X_n)$ to S_ρ, of order*

$$d_{\alpha,p}(S_\rho, S_n) = O\left(\left(\frac{\log n - \log \delta}{n} \right)^{\alpha - \frac{1}{rp}} \right).$$

Corollary 4.1 and remark 3.4 are valid for all n such that $k_n^* \leq n$ (or equivalently such that $n^{r-1}(\log n - \log \delta) \geq q/9$). Note that, because ρ is supported on the unit ball, its covariance has eigenvalues no greater than one, and therefore it must be $q < 1$. It thus suffices to require that $n > 3$ to ensure the condition $k_n^* \leq n$ to hold.

3.1 Proofs

We provide here a proof of our main theoretical result (Theorem 3.1), with some of the intermediate technical results included in the Appendices. For each $\lambda > 0$, we denote by $r^\lambda(x) := \mathbf{1}\{x > \lambda\}$ the step function with a cut-off at λ. Given an empirical covariance operator C_n, we will consider the truncated version $r^\lambda(C_n)$ where, in this notation, r^λ is applied to the eigenvalues of C_n, that is, $r^\lambda(C_n)$ has the same eigen-structure as C_n, but its eigenvalues that are less or equal to λ are clamped to zero.

In order to prove the bound of Equation (3), we begin by proving a more general upper bound of $d_{\alpha,p}(S_\rho, \hat{S}_n^k)$, which is split into a random, and a deterministic part. The bound holds for all values of a free parameter $t > 0$, which is then constrained and optimized in order to find the (close to) tightest version of the bound.

Lemma 3.5. *Let $t > 0$, $0 \leq \alpha \leq \frac{1}{2}$, and $\lambda = \sigma_k(C)$ be the k-th top eigenvalue of C_n, it is,*

$$d_{\alpha,p}(S_\rho, \hat{S}_n^k) \;\leq\; \underbrace{\|(C+tI)^{\frac{1}{2}}(C_n+tI)^{-\frac{1}{2}}\|_\infty^{2\alpha}}_{\mathcal{A}} \;\cdot\; \underbrace{\{3/2(\lambda+t)\}^\alpha}_{\mathcal{B}} \;\cdot\; \underbrace{\|C^\alpha(C+tI)^{-\alpha}\|_p}_{\mathcal{C}} \tag{5}$$

Proof. Let the shorthands $P_\rho := P_{S_\rho}$, and $P_n^\lambda := P_{\hat{S}_n^k}$ denote the metric projection operators onto S_ρ and $\hat{S}_n^k$, respectively. By the definition of $d_{\alpha,p}$ (Eq. (2)), and for all $t > 0$, it is

$$\begin{aligned} d_{\alpha,p}(S_\rho, \hat{S}_n^k) &= \|(P_\rho - P_n^k)C^\alpha\|_p = \|(P_\rho - P_n^k)(C_n+tI)^\alpha(C_n+tI)^{-\alpha}C^\alpha\|_p & (6)\\ &\leq \|(P_\rho - P_n^k)(C_n+tI)^\alpha\|_\infty \;\cdot\; \|(C_n+tI)^{-\alpha}C^\alpha\|_p & (7) \end{aligned}$$

We now bound the two terms of Equation 7.

[**Bound** $\|(P_\rho - P_n^k)(C_n + tI)^\alpha\|_\infty \leq \mathcal{B}$]. Since P_ρ is a metric (orthogonal) projection onto a linear subspace, then clearly it is *dominated* by the identity: $P_\rho \preceq I$, where $\preceq$ is Löwner's partial ordering (Appendix B). Moreover P_ρ and P_n^k are commutative because P_n^k is associated to a linear space that is a subspace of the one associated to P_ρ. Thus applying Lemma B.2 we have $\|(P_\rho - P_n^k)(C_n + tI)^\alpha\|_\infty \leq \|(I - P_n^k)(C_n + tI)^\alpha\|_\infty$. and therefore

$$\begin{aligned} \|(P_\rho - P_n^k)(C_n + tI)^\alpha\|_\infty &\underset{(P_\rho \preceq I)}{\leq} \|(I - P_n^k)(C_n + tI)^\alpha\|_\infty \\ &\leq (\sigma_k(C_n) + t)^\alpha \end{aligned} \tag{8}$$

where the last inequality follows from the bound

$$\|(I - P_n^k)(C_n + tI)^\alpha\|_\infty \leq \sup_{\sigma \in [0,1]} \left(1 - r^\lambda(\sigma)\right)(\sigma + t)^\alpha \leq (\sigma_k(C_n) + t)^\alpha. \tag{9}$$

Note that the middle expression in Equation 9 comes from the function

$$Q_\mu(\alpha, \lambda) := \sup_{0 \leq \sigma \leq 1} (1 - \mu^\lambda(\sigma))\sigma^\alpha$$

of a general regularizing function μ^λ, whose definition is standard in the theory of inverse problems Engl et al. (1996). In the case of the step function r^λ, it is $Q_r(\alpha, \lambda) = \lambda^\alpha$.

The inequality is proven applying Lemma 3.6, indeed we have $\left\|(C + tI)^{\frac{1}{2}}(C_n + tI)^{-\frac{1}{2}}\right\|_\infty^2 \geq \frac{2}{3}$. This is equivalent to $C_n + t \preceq \frac{3}{2}(C + t)$ by (Lemma B.2 point 4), where the $\preceq$ is the Löwner partial order. Note that given two positive semidefinite compact operators A, B, $A \preceq B$ implies that $\sigma_k(A) \leq \sigma_k(B)$ for each $k \geq 1$ (Gohberg et al. (2003) page 186). Thus $\sigma_k(C_n) + t \leq \frac{3}{2}(\sigma_k(C) + t)$ and finally

$$\|(P_\rho - P_n^k)(C_n + tI)^\alpha\|_\infty \leq \{3/2(\lambda + t)\}^\alpha$$

.

[**Bound** $\|(C_n + tI)^{-\alpha} C^\alpha\|_p \leq \mathcal{A} \cdot \mathcal{C}$]. The right-hand side term of Equation 7 can be bounded from above by letting $A_n := (C_n + tI)^{-1}(C + tI)$, and noting that

$$\begin{aligned} \|(C_n + tI)^{-\alpha} C^\alpha\|_p &= \|(C_n + tI)^{-\alpha}(C + tI)^\alpha (C + tI)^{-\alpha} C^\alpha\|_p && (10) \\ &\leq \|(C_n + tI)^{-\alpha}(C + tI)^\alpha\|_\infty \|(C + tI)^{-\alpha} C^\alpha\|_p && (11) \\ &\leq \|(C_n + tI)^{-\frac{1}{2}}(C + tI)^{\frac{1}{2}}\|_\infty^{2\alpha} \|(C + tI)^{-\alpha} C^\alpha\|_p && (12) \\ &= \mathcal{A} \cdot \mathcal{C} && (13) \end{aligned}$$

where all the steps, except for Equation 12, are simple substitutions and rearrangements. As for Equation 12, it holds by Cordes inequality Furuta (1989). □

Note that the right-hand side of Equation (5) is the product of three terms, the left of which ($\mathcal{A}$) involves the empirical covariance operator C_n, which is a random variable, and the right two ($\mathcal{B}$, $\mathcal{C}$) are entirely deterministic. While the term $\mathcal{B}$ has already been reduced to the known quantitiest, α, λ, the remaining terms are bound next. We next bound each term in turn ($\mathcal{B}$ is already We bound the random term $\mathcal{A}$ in the next Lemma, whose proof (postponed to Appendix **??**) makes use of recent concentration results Tropp (2012).

Lemma 3.6 (Term $\mathcal{A}$)**.** *Let* $0 \leq \alpha \leq \frac{1}{2}$*, if* $\frac{9}{n} \log \frac{n}{\delta} \leq t \leq \|C\|_\infty$*. Then with probability* $1 - \delta$ *it is*

$$(2/3)^\alpha \leq \|(C + tI)^{\frac{1}{2}}(C_n + tI)^{-\frac{1}{2}}\|_\infty^{2\alpha} \leq 2^\alpha$$

Proof. By defining the operator $B_n := (C+tI)^{-1/2}(C-C_n)(C+tI)^{-1/2}$, it is simple to verify that it is

$$\|(C+tI)^{\frac{1}{2}}(C_n+tI)^{-\frac{1}{2}}\|_\infty^{2\alpha} = \|(C+tI)^{\frac{1}{2}}(C_n+tI)^{-1}(C+tI)^{\frac{1}{2}}\|_\infty^{\alpha} = \|(I-B_n)^{-1}\|_\infty^{\alpha} \leq (1-\|B_n\|_\infty)^{-\alpha} \quad (14)$$

where the last inequality follows from the fact that $(I-B_n)^{-1} \preceq (1-\|B_n\|_\infty)^{-1} I$ whenever $\|B_n\|_\infty < 1$, and property 3 of Lemma B.2. We now prove a probabilistic upper bound for $\|B_n\|_\infty$.

In order to bound $\|B_n\|_\infty$, we make use of Theorem A.1, which is included in Appendix A for completeness. In particular, we set the parameters of Theorem A.1 as follows. Let $Z := U \otimes U$, with $U := (C+tI)^{-1/2}X$, be a random variable, where $X \sim \rho$ is the random variable from which the data is sampled. Since it is

$$\|Z\|_\infty \leq \|(C+tI)^{-1}\|_\infty \|X\|_{\mathcal{H}}^2 \leq 1/t,$$

we let $R := 1/t$, and $T := \mathbb{E}[Z] = C(C+tI)^{-1}$. Since it is

$$\mathbb{E}_{X\sim\rho}[(U\otimes U - T)^2] = \mathbb{E}_{X\sim\rho}[\|U\|_{\mathcal{H}}^2 U\otimes U - T^2] \preceq \mathbb{E}_{X\sim\rho}[\|U\|_{\mathcal{H}}^2 U\otimes U] \preceq RT,$$

we set $S := RT$. Finally, it is $\sigma^2 = \|RT\|_\infty \leq 1/t$, and $d = \|S\|_1/\|S\|_\infty \leq \frac{(\|C\|_\infty + t)\|T\|_1}{\|C\|_\infty}$. With this choice of parameters, Theorem A.1 implies that, with probability $1-\delta$, it is

$$\|B_n\|_\infty \leq \frac{2\beta}{tn} + \sqrt{\frac{2\beta}{3tn}} \quad (15)$$

with $\beta = \log \frac{4(\|C\|_\infty + t)\|T\|_1}{\|C\|_\infty}$.

By requiring that $t \geq 9\beta/n \geq 4(4+\sqrt{7})\beta/3n$, it can be verified (by substituting $4(4+\sqrt{7})\beta/3n$ into the right-hand side of Equation 15) that this implies $\mathbb{P}[\|B_n\|_\infty \leq 1/2] \geq 1-\delta$. The expression $t \geq 9\beta/n$, however, isn't yet a condition on t, since the right-hand side of the inequality still depends on t. Although we not may solve the inequality for t in closed-form, it is easy to verify that the condition $t \geq \frac{9}{n}\log\frac{n}{\delta}$ is sufficient to ensure that it is satisfied.

Finally, since $t \geq \frac{9}{n}\log\frac{n}{\delta}$ implies $\mathbb{P}[\|B_n\|_\infty \leq 1/2] \geq 1-\delta$ then, with probability $1-\delta$, it holds

$$(2/3)^\alpha \leq (1+\|B_n\|_\infty)^{-\alpha} \leq \|(C+tI)^{\frac{1}{2}}(C_n+tI)^{-\frac{1}{2}}\|_\infty^{2\alpha} \leq (1-\|B_n\|_\infty)^{-\alpha} \leq 2^\alpha$$

as claimed. □

Lemma 3.7 (Term $\mathcal{C}$)**.** *Let C be a symmetric, bounded, positive semidefinite linear operator on $\mathcal{H}$. If $\sigma_k(C) \leq f(k)$ for $k \in \mathbb{N}$, where f is a decreasing function then, for all $t > 0$ and $\alpha \geq 0$, it holds*

$$\left\|C^\alpha (C+tI)^{-\alpha}\right\|_p \leq \inf_{0\leq u\leq 1} g_{u\alpha} t^{-u\alpha} \quad (16)$$

where $g_{u\alpha} = \left(f(1)^{u\alpha p} + \int_1^\infty f(x)^{u\alpha p} dx\right)^{1/p}$. Furthermore, if $f(k) = gk^{-1/\gamma}$, with $0 < \gamma < 1$ and $\alpha p > \gamma$, then it holds

$$\left\|C^\alpha (C+tI)^{-\alpha}\right\|_p \leq Q t^{-\gamma/p} \quad (17)$$

where $Q = (g^\gamma \Gamma(\alpha p - \gamma)\Gamma(1+\gamma)/\Gamma(\gamma))^{1/p}$.

Proof. Since for any $0 \leq u \leq 1$ such that $\|C^{u\alpha}\|_p < \infty$, it is

$$\left\|C^\alpha (C+tI)^{-\alpha}\right\|_p = \left\|C(C+tI)^{-1}\right\|_{\alpha p}^{\alpha} \leq \|C^u\|_{\alpha p}^{\alpha} \left\|C^{1-u}(C+tI)^{-1}\right\|_\infty^{\alpha},$$

and considering that $C^{1-u} \preceq (C+tI)^{1-u}$ (since $C \succeq 0$ and $t > 0$) then, by property 3 of Lemma B.2, it is $\left\|C^{1-u}(C+tI)^{-1}\right\|_\infty \leq \|(C+tI)^{-u}\|_\infty \leq t^{-u}$. Therefore, it follows that

$$\left\|C^\alpha (C+tI)^{-\alpha}\right\|_p \leq \|C^u\|_{\alpha p}^{\alpha} t^{-u\alpha} = \|C^{u\alpha}\|_p t^{-u\alpha}.$$

Since f is decreasing, it follows from the definition of p-Schatten norm ($\|C\|_p^p = \sum_{k\geq 1} \sigma_k(C)^p$) that

$$\|C^{u\alpha}\|_p = \|C\|_{u\alpha p}^{u\alpha} = (\sum_{n\geq 1} \sigma_k(C)^{u\alpha p})^{1/p} \leq \left(f(1)^{u\alpha p} + \int_1^\infty f(x)^{u\alpha p} dx \right)^{1/p}.$$

Given a specific upper bound $f(k)$ of the spectrum of C, we may calculate an upper bound of $\|C^\alpha (C+tI)^{-\alpha}\|_p$. In particular, if $f(k) = gk^{-1/\gamma}$, then $\|C^\alpha (C+tI)^{-\alpha}\|_p^p = \sum_{k\geq 1} h(\sigma_k(C))$ where $h(x) := x^{\alpha p}(x+t)^{-\alpha p}$. Since h is increasing, then $h \circ f$ is decreasing (the composition of increasing and decreasing functions is decreasing), and $h(\sigma_k(C)) \leq h(f(k))$. Therefore $\|C^\alpha (C+tI)^{-\alpha}\|_p^p = \sum_{k\geq 1} h(\sigma_k(C)) \leq \sum_{k\geq 1} h(f(k)) \leq \int_0^\infty h(f(x))dx$, which leads to Equation 17 by plugging in the expression of f and h. □

Proof of Theorem 3.1. The combination of Lemmas 3.5 and 3.6. □

Proof of Theorem 3.2. Application of Lemma 3.7 to Theorem 3.1. □

Finally, Corollary 4.1 is simply a particular case for the reconstruction error $d_R(S_\rho, \cdot) = d_{\alpha,p}(S_\rho, \cdot)^2$, with $\alpha = 1/2, p = 2$.

As noted in Section 3, looser bounds would be obtained if classical Bernstein inequalities in Hilbert spaces Pinelis (1994) were used instead. In particular, Lemma 3.6 would result in a range for t of $qn^{-r/(r+1)} \leq t \leq \|C\|_\infty$, implying $k^* = O(n^{1/(r+1)})$ rather than $O(n^{1/r})$, and thus Theorem 3.2 would become (for $k \geq k^*$) $d_{\alpha,p}(S_\rho, S_n^k) = O(n^{-\alpha r/(r+1)+1/(p(r+1))})$ (compared with the sharper $O(n^{-\alpha+1/rp})$ of Theorem 3.2). For instance, for $p = 2$, $\alpha = 1/2$, and a decay rate $r = 2$ (as in the example of Section 5), it would be: $d_{1/2,2}(S_\rho, S_n) = O(n^{-1/4})$ using Theorem 3.2, and $d_{1/2,2}(S_\rho, S_n) = O(n^{-1/6})$ using classical Bernstein inequalities.

4 Applications of subspace learning

We describe next some of the main uses of subspace learning in the literature.

4.1 Kernel PCA and embedding methods

One of the main applications of subspace learning is in reducing the dimensionality of the input. In particular, one may find nested subspaces of dimension $1 \leq k \leq n$ that minimize the distances from the original to the projected samples. This procedure is known as the Karhunen-Loève, PCA, or Hotelling transform Jolliffe (2005), and has been generalized to reproducing-kernel Hilbert spaces (RKHS) Schölkopf et al. (1997).

In particular, the above procedure amounts to computing an eigen-decomposition of the empirical covariance (Sec. 2.1):

$$C_n = \sum_{i=1}^n \sigma_i u_i \otimes u_i,$$

where the k-th subspace estimate is $\hat{S}_n^k := \operatorname{Ran} C_n^k = \operatorname{span}\{u_i : 1 \leq i \leq k\}$. Note that, in the general case of kernel PCA, we assume the samples $\{x_i\}_{1\leq i\leq n}$ to be in some Reproducing Kernel Hilbert Space (RKHS) $\mathcal{H}$, which are obtained from the observed variables $(z_1, \ldots, z_n) \in Z^n$, for some space Z, through an embedding $x_i := \phi(z_i)$. Typically, due to the very high dimensionality of $\mathcal{H}$, we may only have indirect information about ϕ in the form a kernel function $K : Z \times Z \to \mathbb{R}$: a symmetric, positive definite function satisfying $K(z, w) = \langle \phi(z), \phi(w) \rangle_{\mathcal{H}}$ Steinwart and Christmann (2008) (for technical reasons, we also assume K to be continuous). Note that every such K has a unique associated RKHS, and viceversa Steinwart and Christmann (2008, p. 120–121), whereas, given K, the embedding ϕ is only unique up to an inner product-preserving transformation.

Given a point $z \in Z$, we can make use of K to compute the coordinates of the projection of its embedding $\phi(z)$ onto $\hat{S}_n^k \subseteq \mathcal{H}$ by means of a simple k-truncated eigen-decomposition of K_n, as described in Appendix C.

It is easy to see that the k-truncated kernel PCA subspace $\hat{S}_n^k$ minimizes the empirical reconstruction error $d_R(\hat{S}_n, \hat{S})$, among all subspaces $\hat{S}$ of dimension k. Indeed, it is

$$\begin{aligned} d_R(\hat{S}_n, \hat{S}) &= \mathbb{E}_{x\sim\hat{\rho}} \|x - P_{\hat{S}}(x)\|_{\mathcal{H}}^2 = \mathbb{E}_{x\sim\hat{\rho}} \left\langle (I - P_{\hat{S}})x, (I - P_{\hat{S}})x \right\rangle_{\mathcal{H}} \\ &= \mathbb{E}_{x\sim\hat{\rho}} \left\langle I - P_{\hat{S}}, x \otimes x \right\rangle_{HS} = \left\langle I - P_{\hat{S}}, C_n \right\rangle_{HS}, \end{aligned} \tag{18}$$

where $\langle \cdot, \cdot \rangle_{HS}$ is the Hilbert-Schmidt inner product, form which it is easy to see that the k-dimensional subspace minimizing Equation 18 (alternatively maximizing $< P_{\hat{S}}, C_n >$) is spanned by the k-top eigenvectors of C_n.

Since we are interested in the expected $d_R(S_\rho, \hat{S}_n^k)$ (rather than the empirical $d_R(\hat{S}_n, \hat{S})$) error of the kernel PCA estimate, we may obtain a learning rate for Equation 18 by particularizing Theorem 3.2 to the reconstruction error, for all k (Theorem 3.2), and for $k \geq k^*$ with a suitable choice of k^* (Corollary 4.1). In particular, recalling that $d_R(S_\rho, \cdot) = d_{\alpha,p}(S_\rho, \cdot)^2$ with $\alpha = 1/2$ and $p = 2$, and choosing a value of $k \geq k_n^*$ that minimizes the bound of Theorem 3.2, we obtain the following result.

Corollary 4.1 (Performance of PCA / Reconstruction error)**.** *Let C have eigenvalue decay rate of order r, and $\hat{S}_n^*$ be as in Corollary 3.3. Then it holds, with probability $1 - \delta$,*

$$d_R(S_\rho, \hat{S}_n^*) = O\left(\left(\frac{\log n - \log \delta}{n} \right)^{1-1/r} \right).$$

4.2 Support estimation

The problem of support estimation consists in recovering the support M of a distribution ρ on a metric space Z from identical and independent samples $Z_n = (z_i)_{1 \leq i \leq n}$. We briefly recall a recently proposed approach to support estimation based on subspace learning De Vito et al. (2010), and discuss how our results specialize to this setting, producing a qualitative improvement to theirs.

Given a suitable reproducing kernel K on Z (with associated feature map ϕ), the support M can be characterized in terms of the subspace $S_\rho = \overline{\text{span } \phi(M)} \subseteq \mathcal{H}$ De Vito et al. (2010). More precisely, letting $d_V(x) = \|x - P_V x\|_{\mathcal{H}}$ be the point-subspace distance to a subspace V, it can be shown (see De Vito et al. (2010)) that, if the kernel *separates* [1] M, then it is

$$M = \{z \in Z \mid d_{S_\rho}(\phi(z)) = 0\}.$$

This suggests an empirical estimate $\hat{M} = \{z \in Z \mid d_{\hat{S}}(\phi(z)) \leq \tau\}$ of M, where $\hat{S} = \overline{\text{span } \phi(Z_n)}$, and $\tau > 0$. With this choice, almost sure convergence $\lim_{n\to\infty} d_H(M, \hat{M}) = 0$ in the Hausdorff distance Beer (1993) is related to the convergence of $\hat{S}$ to S_ρ De Vito et al. (2010). More precisely, if the eigenfunctions of the covariance operator $C = \mathbb{E}_{z\sim\rho}[\phi(z) \otimes \phi(z)]$ are uniformly bounded, then it suffices for Hausdorff convergence to bound from above $d_{\frac{r-1}{2r},\infty}$ (where $r > 1$ is the eigenvalue decay rate of C). The following results specializes Corollary 3.3 to this setting.

Corollary 4.2 (Performance of set learning)**.** *If $0 \leq \alpha \leq \frac{1}{2}$, then it holds, with probability $1 - \delta$,*

$$d_{\alpha,\infty}(S_\rho, \hat{S}_n^*) = O\left(\left(\frac{\log n - \log \delta}{n} \right)^{\alpha} \right).$$

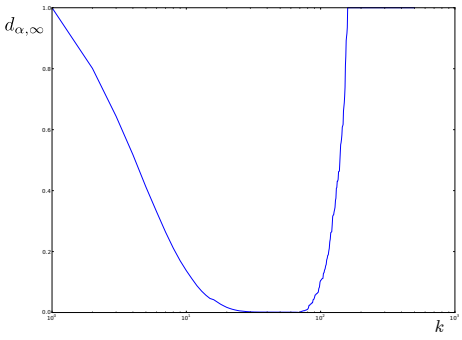

Figure 1: The figure shows the experimental behavior of the distance $d_{\alpha,\infty}(\hat{S}^k, S_\rho)$ between the empirical and the real subspaces, with respect to the regularization parameter. The setting is the one of section 5. Here the real subspace is analytically computed, while the empirical one is computed on a dataset with $n = 1000$ and 32bit floating point precision. Note the numerical instability as k tends to 1000.

Letting $\alpha = \frac{r-1}{2r}$ above yields a high probability bound of order $O\left(n^{-\frac{r-1}{2r}}\right)$ (up to logarithmic factors), which is considerably sharper than the bound $O\left(n^{-\frac{r-1}{2(3r-1)}}\right)$ found in De Vito et al. (2012) (Theorem 7). Note that these are upper bounds for the best possible choice of k (which minimizes the bound). While the optima of both bounds vanish with $n \to \infty$, their behavior is qualitatively different. In particular, the bound of De Vito et al. (2012) is U-shaped, and diverges for $k = n$, while ours is L-shaped (no trade-off), and thus also convergent for $k = n$. Therefore, when compared with De Vito et al. (2012), our results suggest that no regularization is required from a statistical point of view though, as clarified in the following remark, it may be needed for purposes of numerical stability.

Remark 4.3. *While, as proven in Corollary 4.2, regularization is not needed from a statistical perspective, it can play a role in ensuring numerical stability in practice. Indeed, in order to find $\hat{M}$, we compute $d_{\hat{S}}(\phi(z))$ with $z \in Z$. Using the reproducing property of K, it can be shown that, for $z \in Z$, it is $d_{\hat{S}^k}(\phi(z)) = K(z,z) - \left\langle t_z, (\hat{K}_n^k)^\dagger t_z \right\rangle$ where $(t_z)_i = K(z, z_i)$, $\hat{K}_n$ is the Gram matrix $(\hat{K}_n)_{ij} = K(z_i, z_j)$, $\hat{K}_n^k$ is the rank-k approximation of $\hat{K}_n$, and $(\hat{K}_n^k)^\dagger$ is the pseudo-inverse of $\hat{K}_n^k$. The computation of $\hat{M}$ therefore requires a matrix inversion, which is prone to instability for high condition numbers. Figure 1 shows the behavior of the error that results from replacing $\hat{S}$ by its k-truncated approximation $\hat{S}^k$. For large values of k, the small eigenvalues of $\hat{S}$ are used in the inversion, leading to numerical instability.*

5 Experiments

In order to validate our analysis empirically, we consider the following experiment. Let ρ be a uniform one-dimensional distribution in the unit interval. We embed ρ into a reproducing-kernel Hilbert space $\mathcal{H}$ using the exponential of the ℓ_1 distance ($k(u,v) = \exp\{-\|u - v\|_1\}$) as kernel. Given n samples drawn from ρ, we compute its empirical covariance in $\mathcal{H}$ (whose spectrum is plotted in Figure 2 (left)), and truncate its eigen-decomposition to obtain a subspace estimate $\hat{S}_n^k$, as described in Section 2.1, and in Appendix C.

Figure 2 (right) is a box plot of reconstruction error $d_R(S_\rho, \hat{S}_n^k)$ associated with the k-truncated kernel-PCA estimate $\hat{S}_n^k$ (the expected distance in $\mathcal{H}$ of samples to $\hat{S}_n^k$), with $n = 1000$ and varying k. While d_R is computed analytically in this example, and S_ρ is fixed, the estimate $\hat{S}_n^k$ is a random variable, and hence the variability in the graph. Notice from the figure that, as pointed out in Blanchard et al. (2007) and discussed in Section 6, the reconstruction error $d_R(S_\rho, \hat{S}_n^k)$ is always a non-increasing function of k, due to the fact that the kernel-PCA estimates are nested: $\hat{S}_n^k \subset \hat{S}_n^{k'}$ for $k < k'$ (see Section 2.1). The graph is highly concentrated around a curve with a steep intial drop,

[1] A kernel is said to separate M if its associated feature map ϕ satisfies $\phi^{-1}(\overline{\text{span}\ \phi(M)}) = M$ (e.g. the Abel kernel is separating).

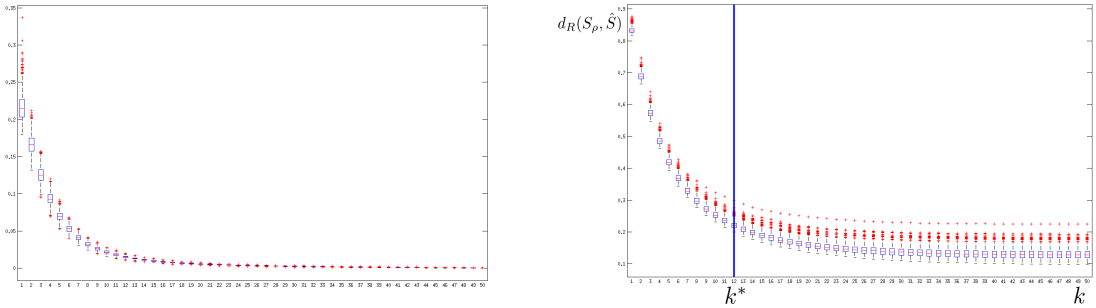

Figure 2: The spectrum of the empirical covariance (left), and the expected distance from a random sample to the empirical k-truncated kernel-PCA subspace estimate (right), as a function of k ($n = 1000$, 1000 trials shown in a boxplot). Our predicted plateau threshold k_n^* (Theorem 3.2) is a good estimate of the value k past which the distance stabilizes.

until reaching some sufficiently high k, past which the reconstruction (pseudo) distance becomes stable, and does not vanish. In our experiments, this behavior is typical for the reconstruction distance and high-dimensional problems.

Due to the simple form of this example, we are able to compute analytically the spectrum of the true covariance C. In this case, the eigenvalues of C decay as $2\gamma/((k\pi)^2 + \gamma^2)$, with $k \in \mathbb{N}$, and therefore they have a polynomial decay rate $r = 2$ (see Section 3). Given the known spectrum decay rate, we can estimate the plateau threshold $k = k_n^*$ in the bound of Theorem 3.2, which can be seen to be a good approximation of the observed start of a plateau in $d_R(S_\rho, \hat{S}_n^k)$ (Figure 2, right). Notice that our bound for this case (Corollary 4.1) similarly predicts a steep performance drop until the threshold $k = k_n^*$ (indicated in the figure by the vertical blue line), and a plateau afterwards.

6 Discussion

Figure 3 shows a comparison of our learning rates with existing rates in the literature Blanchard et al. (2007); Shawe-Taylor et al. (2005). The plot shows the polynomial decay rate c of the bound $d_R(S_\rho, \hat{S}_n^k) = O(n^{-c})$, as a function of the eigenvalue decay rate r of the covariance C, computed at the best value k_n^* (which minimizes the bound).

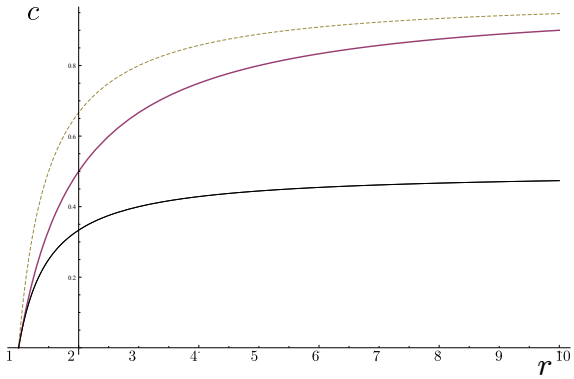

Figure 3: Known upper bounds for the polynomial decay rate c (for the best choice of k), for the expected distance from a random sample to the empirical k-truncated kernel-PCA estimate, as a function of the covariance eigenvalue decay rate (higher is better). Our bound (purple line), consistently outperforms previous ones (Shawe-Taylor et al. (2005) black line). The top (Blanchard et al. (2007) dashed) line, has significantly stronger assumptions, and is only included for completeness.

The rate exponent c, under a polynomial eigenvalue decay assumption for C, is $c = \frac{s(r-1)}{r-s+sr}$ for Blanchard et al. (2007) and $c = \frac{r-1}{2r-1}$ for Shawe-Taylor et al. (2005), where s is related to the fourth moment. Note that, among the two (purple and black) that operate under the same assumptions, our bound (purple line) is the best by a wide margin. The top, best performing, dashed line Blanchard et al. (2007) is obtained for the best possible fourth-order moment constraint $s = 2r$, and is therefore not a fair comparison. However, it is worth noting that our bounds perform almost as well as the most restrictive one, even when we do not include any fourth-order moment constraints.

Choice of truncation parameter k. Since, as pointed out in Section 2.1, the subspace estimates $\hat{S}_n^k$ are nested for increasing k (i.e. $\hat{S}_n^k \subseteq \hat{S}_n^{k'}$ for $k < k'$), the distance $d_{\alpha,p}(S_\rho, \hat{S}_n^k)$, and in particular the reconstruction error $d_R(S_\rho, \hat{S}_n^k)$, is a non-increasing function of k. As has been previously discussed Blanchard et al. (2007), this suggests that there is no tradeoff in the choice of k. Indeed, the fact that the estimates $\hat{S}_n^k$ become increasing close to S_ρ as k increases indicates that the best choice is the highest: $k = n$.

Interestingly, however, both in practice (Section 5), and in theory (Section 3), we observe that a typical behavior for the subspace learning problem in high dimensions (e.g. kernel PCA) is that there is a certain value of $k = k_n^*$, past which performance plateaus. For problems such as spectral embedding methods Tenenbaum et al. (2000); Donoho and Grimes (2003); Weinberger and Saul (2006), in which a degree of dimensionality reduction is desirable, producing an estimate $\hat{S}_n^k$ where k is close to the plateau threshold may be a natural parameter choice: it leads to an estimate of the lowest dimension ($k = k_n^*$), whose distance to the true S_ρ is almost as low as the best-performing one ($k = n$).

References

Tsuyoshi Ando and Xingzhi Zhan. Norm inequalities related to operator monotone functions. *Mathematische Annalen*, 315:771–780, 1999. ISSN 0025-5831. URL `http://dx.doi.org/10.1007/s002080050335`. 10.1007/s002080050335.

G. Beer. *Topologies on Closed and Closed Convex Sets*. Springer, 1993. ISBN 9780792325314.

M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.

Y. Bengio, O. Delalleau, N.L. Roux, J.F. Paiement, P. Vincent, and M. Ouimet. Learning eigenfunctions links spectral embedding and kernel pca. *Neural Computation*, 16(10):2197–2219, 2004a.

Y. Bengio, J.F. Paiement, and al. Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering. *Advances in neural information processing systems*, 16:177–184, 2004b.

S. Bernstein. *The Theory of Probabilities*. Gastehizdat Publishing House, Moscow, 1946.

G. Blanchard, O. Bousquet, and L. Zwald. Statistical properties of kernel principal component analysis. *Machine Learning*, 66(2):259–294, 2007.

I. Borg and P.J.F. Groenen. *Modern multidimensional scaling: Theory and applications*. Springer, 2005.

J.C. Bourin. Some inequalities for norms on matrices and operators. *Linear Algebra and its Applications*, 292 (1):139–154, 1999.

Ernesto De Vito, Lorenzo Rosasco, and Alessandro Toigo. Spectral regularization for support estimation. *Advances in Neural Information Processing Systems, NIPS Foundation*, pages 1–9, 2010.

Ernesto De Vito, Lorenzo Rosasco, and al. Learning sets with separating kernels. *arXiv:1204.3573*, 2012.

D.L. Donoho and C. Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10):5591–5596, 2003.

H.W. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*. Mathematics and Its Applications. Springer, 1996. ISBN 9780792341574. URL http://books.google.de/books?id=2bzgmMv5EVcC.

Takayuki Furuta. Norm inequalities equivalent löwner-heinz theorem. *Reviews in Mathematical Physics*, 01(01): 135–137, 1989.

Israel Gohberg, Seymour Goldberg, and al. *Basic classes of linear operators*. Springer, 2003.

J. Ham, D.D. Lee, S. Mika, and B. Schölkopf. A kernel view of the dimensionality reduction of manifolds. In *Proceedings of the twenty-first international conference on Machine learning*, page 47. ACM, 2004.

F. Hansen. An operator inequality. *Mathematische Annalen*, 246(3):249–250, 1980.

I. Jolliffe. *Principal component analysis*. Wiley Online Library, 2005.

K. T. Löwner. über monotone matrixfunktionen. *Mathematische Zeitschrift*, 38:177–216, 1934.

Andreas Maurer and Massimiliano Pontil. K–dimensional coding schemes in hilbert spaces. *IEEE Transactions on Information Theory*, 56(11):5839–5846, 2010.

Iosif Pinelis. Optimum bounds for the distributions of martingales in banach spaces. *The Annals of Probability*, pages 1679–1706, 1994.

J.R. Retherford. *Hilbert Space: Compact Operators and the Trace Theorem*. London Mathematical Society Student Texts. Cambridge University Press, 1993. ISBN 9780521429337. URL http://books.google.de/books?id=IEixfs1Q514C.

S.T. Roweis and L.K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500): 2323–2326, 2000.

Alessandro Rudi, Guillermo D Canas, and Lorenzo Rosasco. On the Sample Complexity of Subspace Learning. In *Advances in Neural Information Processing Systems*, pages 2067–2075, 2013.

L.K. Saul and S.T. Roweis. Think globally, fit locally: unsupervised learning of low dimensional manifolds. *The Journal of Machine Learning Research*, 4:119–155, 2003.

B. Schölkopf, A. Smola, and K.R. Müller. Kernel principal component analysis. *Artificial Neural Networks-ICANN'97*, pages 583–588, 1997.

J. Shawe-Taylor, C. K. Williams, N. Cristianini, and J. Kandola. On the eigenspectrum of the gram matrix and the generalization error of kernel-pca. *Information Theory, IEEE Transactions on*, 51(7), 2005.

I. Steinwart and A. Christmann. *Support vector machines*. Information science and statistics. Springer-Verlag. New York, 2008. ISBN 9780387772424. URL http://books.google.de/books?id=HUnqnrpYt4IC.

J. Sun, S. Boyd, L. Xiao, and P. Diaconis. The fastest mixing markov process on a graph and a connection to a maximum variance unfolding problem. *SIAM review*, 48(4):681–699, 2006.

J.B. Tenenbaum, V. De Silva, and J.C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

J.A. Tropp. User-friendly tools for random matrices: An introduction. 2012.

K.Q. Weinberger and L.K. Saul. Unsupervised learning of image manifolds by semidefinite programming. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 2, pages II–988. IEEE, 2004.

K.Q. Weinberger and L.K. Saul. Unsupervised learning of image manifolds by semidefinite programming. *International Journal of Computer Vision*, 70(1):77–90, 2006.

C.K.I. Williams. On a connection between kernel pca and metric multidimensional scaling. *Machine Learning*, 46(1):11–19, 2002.

A Concentration bounds on compact operators

Theorem A.1. [Tropp's concentration inequality Tropp (2012, Theorem 7.3.1) on the operator norm] *Let $(Z_i)_{1\le i\le n}$ be independent copies of the random variable Z with values in the space of bounded self-adjoint operators $\mathcal{B}(\mathcal{H})$ over a separable Hilbert space $\mathcal{H}$. Define $T := \mathbb{E}[Z]$, and let there be $S \in \mathcal{S}(\mathcal{H})$ such that $\mathbb{E}\left[(Z-T)^2\right] \le S$, and a finite number R such that $\|Z\|_\infty \le R$ almost everywhere. Define the quantities $d := \|S\|_1 / \|S\|_\infty$ and $\sigma^2 := \|S\|_\infty$. Then, for $0 < \delta \le d$, it holds*

$$\mathbb{P}\left\{\left\|\frac{1}{n}\sum_{i=1}^n Z_i - T\right\|_\infty \le \frac{\beta R}{n} + \sqrt{\frac{3\beta\sigma^2}{n}}\right\} \le 1-\delta$$

where $\beta := \frac{2}{3}\log\frac{4d}{\delta}$.

Theorem A.2. [Pinelis inequality Pinelis (1994) on the Hilbert-Schmidt norm] *Let $(Z_i)_{1\le i\le n}$ be independent copies of the random variable Z with values in the space of bounded operators $\mathcal{B}(\mathcal{H})$ over a separable Hilbert space $\mathcal{H}$. Define $T := \mathbb{E}[Z]$, and let there be $S \in \mathcal{S}(\mathcal{H})$ such that $\mathbb{E}\left[(Z-T)^2\right] \le S$, and a finite number R such that $\|Z\|_2 \le R$ almost everywhere. Define the quantity $\sigma^2 := \operatorname{Tr} S$. Then, for $\delta > 0$, it holds*

$$\mathbb{P}\left\{\left\|\frac{1}{n}\sum_{i=1}^n Z_i - T\right\|_2 \le \frac{\beta R}{n} + \sqrt{\frac{3\beta\sigma^2}{n}}\right\} \le 1-\delta$$

where $\beta := \frac{2}{3}\log\frac{1}{\delta}$.

Note that this theorem corresponds to the classical Bernstein inequality Bernstein (1946) when $\mathcal{H} = \mathbb{R}$.

B Properties of positive semidefinite operators

Let $\mathcal{B}(\mathcal{H})$ be the space of bounded linear operators in $\mathcal{H}$.

Definition B.1 (Löwner's partial ordering)**.** *Given positive semidefinite operators $A, B \in \mathcal{B}(\mathcal{H})$, it is $A \preceq B$ if $B - A$ is positive semidefinite (i.e. $\langle f, Af\rangle \le \langle f, Bf\rangle$, for all $f \in \mathcal{H}$).*

Lemma B.2 (Properties of Löwner's partial ordering Bourin (1999); Ando and Zhan (1999))**.** *Let $A, B \in \mathcal{B}(\mathcal{H})$ be positive semidefinite such that $A \preceq B$, and $C, D, E \in \mathcal{B}(\mathcal{H})$, $0 \le r \le 1$ and $\|\cdot\|_p$ the p-Schatten norm with $p \ge 0$, then*

1. *$CAC^* \preceq CBC^*$*
2. *$A^r \preceq B^r$ and $CA^rC^* \preceq \|C\|_\infty^{2-2r}(CAC^*)^r$*
3. *$\|A\|_p \le \|B\|_p$ and $\|DC\|_p \le \|EC\|_p$ whenever $D^*D \preceq E^*E$.*

Proof. 1) For all $x \in \mathcal{H}$, and renaming $y := C^*x$, it is $\langle x, CAC^*x\rangle = \langle y, Ay\rangle \le \langle y, By\rangle = \langle x, CBC^*x\rangle$.

2) $A^r \preceq B^r$ is the well-known Löwner inequality Löwner (1934), and by Hansen (1980), it is $WA^rW^* \preceq (WA^rW^*)^r$ when $\|W\|_\infty \le 1$.

3) When $p \ge 1$, it is $\operatorname{Tr}(U+V)^p \ge \operatorname{Tr}U^p + \operatorname{Tr}V^t$ when $U, V \ge 0$ (see Bourin (1999); Ando and Zhan (1999) and references therein). Therefore, it holds

$$\|B\|_p^p = \|A + (B-A)\|_p^p = \operatorname{Tr}(A+(B-A))^p \ge \operatorname{Tr}A^p + \operatorname{Tr}(B-A)^p \ge \operatorname{Tr}A^p = \|A\|_p^p$$

When $0 \le p \le 1$, the fact that $A \preceq B$ implies $A^p \preceq B^p$, and therefore, by the preceding argument it is $\|A\|_p^{\frac{1}{p}} = \|A^p\|_1 \le \|B^p\|_1 = \|B\|_p^{\frac{1}{p}}$.

Finally since it is, by definition, $\|D\|_p = \|D^*D\|_{\frac{p}{2}}^{\frac{1}{2}}$ then, by the assumption $D^*D \preceq E^*E$, we have that $\|DC\|_p^2 = \|C^*D^*DC\|_{\frac{p}{2}} \le \|C^*E^*EC\|_{\frac{p}{2}} = \|EC\|_p^2$. □

C Auxiliary proofs

Let ρ be a probability measure supported in the unit ball of a separable Hilbert space $\mathcal{H}$, with an associated covariance operator C, and $\Lambda_\rho(\mathcal{H})$ be the set of linear subspaces contained in the span S_ρ of the support of ρ. For each $U \in \Lambda_\rho(\mathcal{H})$, let $\|U\|_{\alpha,p} := \|P_U C^\alpha\|_p$, where P_U is the orthogonal projection operator onto U.

Proposition C.1. *$(\Lambda_\rho(\mathcal{H}), \|\cdot\|_{\alpha,p})$ with $0 \leq \alpha \leq 1$, $1 \leq p \leq \infty$ is a Banach space.*

Proof. From the definition of $\|\cdot\|_{\alpha,p}$, it is clear that all the norm properties are direct except for identifiability. Let $U \in \Lambda_\rho(\mathcal{H})$, and therefore $U \subseteq S_\rho$. Since $S_\rho = \operatorname{Ran} C$, then clearly $S_\rho \cap \ker C^\alpha = \emptyset$, and this is true even for $\alpha = 0$. Let $\|U\|_{\alpha,p} = \|P_U C^\alpha\|_p = 0$. Since $\|\cdot\|_p$ is a norm, it must be $P_U C^\alpha = 0$, or equivalently $U \subseteq \ker C^\alpha$. Since $U \subseteq S_\rho$ and $S_\rho \cap \ker C^\alpha = \emptyset$, it is $P_U = 0$, and therefore $U = \{0\}$. □

Proposition C.2. *Given ρ a Borel probability measure with support in the unit ball of a separable Hilbert space $\mathcal{H}$, its second order moment $C = \mathbb{E}_{X\sim\rho} X \otimes X$ is a symmetric, positive semidefinite, compact linear operator with $\|C\|_1 \leq 1$.*

Proof. C is symmetric by virtue of being a sum of symmetric terms. For $u \in \mathcal{H}$, it is $\langle u, Cu\rangle = \int_{\mathcal{H}} \langle x, u\rangle \, d\rho(x) \geq 0$. Finally, to prove that C is compact, we show that its 1-norm is finite:

$$\|C\|_1 = \operatorname{Tr}\left(\int_{\mathcal{H}} x \otimes x d\rho(x)\right) = \int_{\mathcal{H}} \operatorname{Tr}(x \otimes x) d\rho(x) = \int_{\mathcal{H}} \langle x, x\rangle_{\mathcal{H}} \, d\rho(x) \leq 1.$$

□

Proposition C.3. *The span of the support of ρ is the range of its covariance operator: $S_\rho = \overline{\operatorname{Ran} C}$.*

Proof. Since C is self-ajoint, then $\operatorname{Ran} C$ and $\ker C$ are orthogonal complements Retherford (1993, p. 58), and therefore $S_\rho = \overline{\operatorname{Ran} C}$ is equivalent to $S_\rho^\perp = \ker C$.
$[S_\rho^\perp \supseteq \ker C]$. By the definition of C it is clear that every $u \in \ker C$ is orthogonal to S_ρ. Furthermore, any $v \in S_\rho$ can be written as an infinite linear combination of vectors in the support of ρ: $v = \lim_{k\to\infty} \sum_{i=1}^k \lambda_i x_i$, where $x_i \in \operatorname{supp}\rho$. Since the dot product is continuous, for all $u \in \ker C$, it holds

$$\langle u, v\rangle_{\mathcal{H}} = \left\langle u, \lim_{k\to\infty} \sum_{i=1}^k \lambda_i x_i \right\rangle_{\mathcal{H}} = \lim_{k\to\infty} \sum_{i=1}^k \lambda_i \langle u, x_i\rangle_{\mathcal{H}} = 0$$

and therefore $\ker C \subseteq S_\rho^\perp$.
$[S_\rho^\perp \subseteq \ker C]$. Let $u \in S_\rho^\perp$, and therefore, in particular u is orthogonal to every vector in $\operatorname{supp}\rho$. If $v \in \operatorname{Ran} C$ then, by definition, there is $w \in \mathcal{H}$ such that $v = Cw$. Finally, it is

$$\langle u, v\rangle_{\mathcal{H}} = \langle u, Cw\rangle_{\mathcal{H}} = \left\langle u, \int_{\mathcal{H}} x \langle x, w\rangle_{\mathcal{H}} \, d\rho(x)\right\rangle = \int_{\mathcal{H}} \langle u, x\rangle_{\mathcal{H}} \langle x, w\rangle_{\mathcal{H}} \, d\rho(x) \underset{\langle u,x\rangle=0}{=} 0$$

□

Kernel principal component analysis. Let $K : Z \times Z \to \mathbb{R}$ be a continuous kernel (a symmeric, positive definite function), on a space Z, and ϕ (one of) its corresponding feature map onto a reproducing-kernel Hilbert space $\mathcal{H}$, such that $K(u, v) = \langle \phi(u), \phi(v)\rangle_{\mathcal{H}}$. Given samples $(z_i)_{1\leq i\leq n}$ in Z, let $K_n \in \mathbb{R}^{n\times n}$ be the symmetric, positive semidefinite matrix with entries $(K_n)_{ij} = K(z_i, z_j)$, and $K_n = V\Sigma V^*$ its eigen-decomposition, with eigenvalues $\{\sigma_i\}_{1\leq i\leq n}$ and eigenvectors $\{v_i\}_{1\leq i\leq n}$. Letting $S^* = [\phi(z_1)\dots\phi(z_n)] = [x_1\dots x_n]$ be the embedded samples, then it is $K_n = SS^*$, and $C_n = \frac{1}{n}S^*S$,

and therefore K_n and C_n have the same eigenvalues (up to a factor of n). By considering the eigen-decomposition $C_n = \frac{1}{n} U \Sigma U^*$, where $U : \mathbb{R}^n \to \mathcal{H}$ is $U = [u_1 \dots u_n]$, it follows that $U = S^* V \Sigma^{-1/2}$ and therefore, for $z \in Z$, the projection of $\phi(z)$ onto the j-th top eigenvector u_j of C_n is

$$\langle \phi(z), u_j \rangle_{\mathcal{H}} = \left\langle \phi(z), \sigma_j^{-1/2} \sum_{l=1}^{n} x_l v_{jl} \right\rangle_{\mathcal{H}} = \sigma_j^{-1/2} \sum_{l=1}^{n} K(x, k_l) v_{jl} \tag{19}$$

where v_{jl} is the l-th coordinate of the j-th eigenvector of K_n. Note that Eq. (19) can be computed, for $1 \leq j \leq k$, from just the k top eigenvectors and eigenvalues of K_n and, assuming K_n given, therefore has the same computational cost as a (k-truncated) $n \times n$ eigen-decomposition.

Proposition C.4. *Let $d_R(S_\rho, \hat{S}) = \mathbb{E}_{x \sim \rho} \|x - P_{\hat{S}}(x)\|_{\mathcal{H}}^2$ be the expected (squared) distance from samples to their projection onto a linear subspace $\hat{S}$, and $d_{\alpha,p}(S_\rho, \hat{S}) = \|(P_{S_\rho} - P_{\hat{S}}) C^\alpha\|_p$. It is $d_R(S_\rho, \hat{S}) = d_{1/2,2}(S_\rho, \hat{S})^2$.*

Proof. By the linearity of the trace, it holds:

$$\begin{aligned}
d_R(S_\rho, \hat{S}) &= \int_{\mathcal{H}} \|x - P_{\hat{S}}(x)\|_{\mathcal{H}}^2 d\rho(x) = \int_{\mathcal{H}} \|(P_{S_\rho}) - P_{\hat{S}})x\|_{\mathcal{H}}^2 d\rho(x) \\
&= \int_{\mathcal{H}} \left\langle x, (P_{S_\rho} - P_{\hat{S}})^2 x \right\rangle_{\mathcal{H}} d\rho(x) = \int_{\mathcal{H}} \operatorname{Tr}\left((P_{S_\rho} - P_{\hat{S}})^2 x \otimes x \right) d\rho(x) \\
&= \operatorname{Tr}\left(P_{S_\rho} - P_{\hat{S}})^2 \int_{\mathcal{H}} x \otimes x d\rho(x) \right) = \operatorname{Tr}\left(C^{1/2} (P_{S_\rho} - P_{\hat{S}})^2 C^{1/2} \right) \\
&= \|(P_{S_\rho} - P_{\hat{S}}) C^{1/2}\|_2^2 = d_{1/2,2}(S_\rho, \hat{S})^2
\end{aligned}$$

□